THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Spider: Deep Learning-driven Sparse Mobile Traffic Measurement Collection and Reconstruction

Citation for published version:

Fang, Y, Diallo, A, Zhang, C & Patras, P 2022, Spider: Deep Learning-driven Sparse Mobile Traffic Measurement Collection and Reconstruction. in *IEEE Global Communications Conference (GLOBECOM)*. IEEE Conference and Exhibition on Global Telecommunications (GLOBECOM), Institute of Electrical and Electronics Engineers (IEEE), 2021 IEEE Global Communications Conference, Madrid, Spain, 7/12/21. <https://doi.org/10.1109/GLOBECOM46510.2021.9685804>

Digital Object Identifier (DOI):

[10.1109/GLOBECOM46510.2021.9685804](https://doi.org/10.1109/GLOBECOM46510.2021.9685804)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Global Communications Conference (GLOBECOM)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Spider: Deep Learning-driven Sparse Mobile Traffic Measurement Collection and Reconstruction

Yini Fang

School of Informatics
University of Edinburgh, UK
yini.fang@ed.ac.uk

Alec F. Diallo

School of Informatics
University of Edinburgh, UK
alec.frenn@ed.ac.uk

Chaoyun Zhang

Lightspeed & Quantum Studios
Tencent, China
vyokkyzhang@tencent.com

Paul Patras

School of Informatics
University of Edinburgh, UK
paul.patras@ed.ac.uk

Abstract—Data-driven mobile network management hinges on accurate traffic measurements, which routinely require expensive specialized equipment and substantial local storage capabilities, and bear high data transfer overheads. To overcome these challenges, in this paper we propose Spider, a deep-learning-driven mobile traffic measurement collection and reconstruction framework, which reduces the cost of data collection while retaining state-of-the-art accuracy in inferring mobile traffic consumption with fine geographic granularity. Spider harnesses Reinforcement Learning and tackles large action spaces to train a policy network that selectively samples a minimal number of cells where data should be collected. We further introduce a fast and accurate neural model that extracts spatiotemporal correlations from historical data to reconstruct network-wide traffic consumption based on sparse measurements. Experiments we conduct with a real-world mobile traffic dataset demonstrate that Spider samples 48% fewer cells as compared to several benchmarks considered, and yields up to 67% lower reconstruction errors than state-of-the-art interpolation methods. Moreover, our framework can adapt to previously unseen traffic patterns.

I. INTRODUCTION

Mobile traffic consumption continues to grow sharply and, as 5G is rolled out, the volume of data traffic per month globally is expected to hit a staggering 226 exabytes milestone by 2026 [1]. In this context, network visibility is becoming critical to the management of resources and to assuring network performance. Data-driven network management hinges on accurate traffic measurements [2] collected by dedicated probes that are deployed, e.g., at Packet Gateways (PGWs) [3]. Yet processing vast amounts of data in a scalable and timely fashion is ever more challenging, as it involves substantial local storage, heavy overhead in transferring detailed logs to central locations for analysis, data filtering by scope (e.g., cell ID, session start/end time, traffic volume/type, etc.) to serve specific use cases [4], and it relies on high-performance computing platforms to extract essential insights. Therefore mobile operators urgently need more cost-effective alternative solutions for traffic monitoring and analysis.

Adaptive sampling and compressed sensing have been employed to simplify the network traffic characterization task, whereby only a subset of data collection points are activated and/or monitoring is performed during sparse intervals. Complete network traffic snapshots are subsequently reconstructed via interpolation [5]–[7]. Although such methods adjust the sampling frequency based on previously measured network

activity, sampling locations are routinely selected at random, which (i) overlooks important spatiotemporal correlations specific to mobile traffic, leading to inaccurate reconstruction; and (ii) such myopic view leaves room for sampling overhead reduction, without sacrificing interpolation fidelity.

More recently, neural network models have been proposed to predict data traffic demand in large coverage areas when only incomplete measurements are available [8], or to perform mobile traffic super-resolution, i.e., infer traffic consumption with fine geographic granularity, based only on coarse aggregate measurements [9]. Although these approaches obtain good-quality interpolations by harnessing the exceptional abstract feature extraction abilities of deep learning, *they are built on the premise that the measurement collection points are fixed*. In practice, mobile operators need intelligent strategies to instrument dynamic measurement collection campaigns using virtual probe modules that can be instantiated as needed.

In this paper, we propose Spider, a deep learning-driven mobile traffic measurement collection and reconstruction framework for infrastructure-level data. Spider relies on a dedicated neural network that we train to selectively sample small subsets of target mobile coverage areas. It then employs a purpose-built mobile traffic reconstruction neural model, which exploits spatiotemporal correlations in historical data, to infer mobile traffic consumption across the entire deployment. Our framework reduces dramatically the cost of data collection while retaining high traffic consumption inference accuracy. In summary, we make the following **key contributions**:

- (1) We introduce a policy network that takes as input sparse historical traffic snapshots and outputs cell selection matrices indicating at which locations to activate measurement collection, so as to minimize overhead while acquiring enough data to ensure high-quality traffic consumption interpolations at all non-sampled cells. We take a Deep Reinforcement Learning (DRL) approach to produce examples for training this policy network. Given the large action space, our DRL agent learns in a tractable manner by sampling small subsets of the action space based on the most likely action and its nearest neighbors. The highest valued action from such subsets is then selected, circumventing the need to evaluate the entire action space.
- (2) We propose MTRNet, a deep neural network specifically tailored to mobile traffic reconstruction from sparse traffic

measurements. MTRNet outperforms existing methods in terms of accuracy (up to 67% lower MAE) and reconstruction speed (up to 835 $\times$ runtime reduction).

- (3) We evaluate our Spider framework using a real-world mobile traffic dataset collected by a major European operator, showing that our solution learns to effectively reconstruct complete traffic matrices from sparse samples, even when applied to previously unseen traffic patterns such as those observed during holidays. On average, Spider accurately reconstructs city-scale traffic snapshots with up to 48% fewer samples than a range of benchmarks considered.

II. PROBLEM FORMULATION

Consider a mobile network coverage area that is geographically partitioned into a grid with $X \times Y$ squares (cells), where a square denotes an atomic region for mobile traffic collection. We denote by F_t the traffic consumption snapshot across all cells, observed over an interval $[t - \Delta, t]$, i.e.,

$$F_t = \begin{bmatrix} d_{1,1}^t & d_{1,2}^t & \dots & d_{1,Y}^t \\ d_{2,1}^t & d_{2,2}^t & \dots & d_{2,Y}^t \\ \vdots & \vdots & \ddots & \vdots \\ d_{Y,1}^t & d_{Y,2}^t & \dots & d_{X,Y}^t \end{bmatrix},$$

where $d_{i,j}^t$ is the volume of traffic at cell (i, j) , and Δ is the time granularity configurable by a network administrator, with which traffic measurement equipment (probes) compute summaries of the volume of traffic observed at different locations. A mobile network operator will need a sampling strategy that gives a binary selection matrix B_t indicating which cells are to be selected for measurement collection at time t , i.e., $b_{i,j}^t = 1$, if cell (i, j) is selected; $b_{i,j}^t = 0$, otherwise; and subsequently use the measurements collected to infer the traffic consumption across the entire deployment. A selection matrix produces a sparse measurement matrix M_t of a network traffic snapshot F_t , i.e.

$$M_t = F_t \circ B_t,$$

where $\circ$ is the Hadamard product. Let $\mathbb{M}_t = [M_{t-T}, \dots, M_t]$ denote the sparse measurement matrices collected over T recent timestamps, and $\hat{F}_t$ the reconstruction of F_t given $\mathbb{M}_t$.

The problem we seek to solve is thus twofold: (i) finding a policy $g(\cdot)$ that outputs cell selection matrices $\tilde{B}_t = g(\mathbb{M}_t)$, which contain minimum numbers of elements $b_{i,j}^t$ set to 1, so that the error of reconstructing F_t using those measurements collected is below a predefined threshold ϵ , by using (ii) a data interpolation algorithm $f(\cdot)$. Formally,

$$\tilde{B}_t := \arg \min_{B_t} \sum_{i,j} b_{i,j}^t, \quad (1)$$

$$\text{s.t. } \text{MAE}(\hat{F}_t, F_t) < \epsilon, \quad (2)$$

$$\hat{F}_t = f(\mathbb{M}_t). \quad (3)$$

In the above MAE is the Mean Absolute Error, i.e.,

$$\text{MAE}(\hat{F}_t, F_t) = \frac{1}{X \times Y} \sum_{i,j} |\hat{d}_{i,j}^t - d_{i,j}^t|. \quad (4)$$

We solve this problem using a set of purpose-built deep neural network models as we explain next.

Fig. 1: Proposed Spider framework: sparse mobile traffic snapshots used by a policy network to select optimal cells where to collect measurements; a dedicated reconstruction neural model outputs completed network traffic snapshot.

III. MOBILE TRAFFIC MEASUREMENT COLLECTION AND RECONSTRUCTION FRAMEWORK

To solve the problem defined by (1)–(3), we propose Spider, an original mobile network traffic measurement collection and reconstruction framework that (i) predicts a small number of optimal locations where traffic should be sampled, and (ii) learns to infer the volume of traffic at locations not selected, so as to reconstruct complete snapshots of the traffic consumption across an entire network deployment. We give a diagrammatic view of our approach in Fig. 1. We design original neural models that harness the unique characteristics of mobile traffic to solve both of these tasks.

A. RL for Cell Selection with Large Action Spaces

To collect measurements with minimum sampling overhead, we first leverage Reinforcement Learning (RL) and train an agent that learns the likelihood of selecting a cell where to sample traffic, based on past experience. We subsequently use this agent to train a policy network that *directly outputs optimal cell selection matrices*, given prior observations, and thus is suitable for real-time decision making.

We start by regarding cell selection as an episodic task, which can be modelled as a Markov Decision Process (MDP), $\mathcal{M} := (S, A, P, r, \gamma)$, where

- S is the set of states s_i , in our case a state representing a collection of sparse measurement matrices of the past T network traffic snapshots, i.e., $s_i = \mathbb{M}_t^i = [M_{t-T}, \dots, M_{t-1}^i]$, with M_{t-1}^i denoting the sparse measurement matrix at iteration $i - 1$ of an episode t ;
- A is the set of possible actions, where an action a_i corresponds to sampling one of all the previously unselected cells at timestamp t , from which the agent can choose;
- $P(s, a, s')$ is the probability that action a in state s will lead to next state s' ;
- $r(s, a, s')$ is the reward the agent receives as a consequence of choosing action a when in state s ; we work with $r(s, a, s') = -\text{MAE}(f(s'), F_t)$ to incentivize the agent to take actions that reduce the reconstruction error; while a range of methods $f(\cdot)$ can be employed for reconstruction, including compressive sensing or K-Nearest Neighbour-based interpolation, our Spider framework evaluates how good an action is using a pre-

Fig. 2: Structure of the DRL agent. The pseudo action given by the neural network is used to generate a set of potential candidates that, when evaluated by MTRNet, allows the agent to select the action that minimizes the reconstruction error.

trained Mobile Traffic Reconstruction neural Network (MTRNet), which we detail in Sec. III-B.

- $\gamma \in [0, 1]$ is the discount rate – an agent’s objective is to maximize a cumulative reward it receives (expected return), i.e. the sum of discounted rewards $r\gamma^{k-1}$ over k iterations; we work with $\gamma = 0$, meaning that at each iteration we seek to maximize the immediate return.

We consider a cell selection episode t to be completed at an iteration I when the inference error is less than a predefined reconstruction quality threshold.

With traditional DRL, finding a policy that maps from states to probabilities of choosing an action in a large action space such as city-scale mobile network deployments comprising thousands of cells would require massive amounts of memory. Further, exploring the action space exhaustively, until sufficient cells are selected to yield usable reconstructed traffic snapshots, would involve impractically large runtimes. To address these issues, we design a DRL agent based on a neural network architecture to model the policy function, which we train using an efficient algorithm that only evaluates a subset of potentially good actions from the whole action space, among which the best action is selected following evaluation. As such the agent will first output a pseudo action $\hat{a}$, which is used to generate a subset of k potential actions by finding the K -Nearest Neighbors of $\hat{a}$ in A based on ℓ_2 distance, i.e.,

$$A'_k = knn(\hat{a}) = \arg \min_{a \in A}^k \|a - \hat{a}\|_2,$$

because the spatial correlation between two cells is related to their distance. To explore previously unseen actions, the generated subset A'_k is expanded with η randomly selected actions $rand_\eta(A'_k)$ not already in A'_k , where η is given by an exponentially decaying function, i.e.

$$A' := A'_{k-\eta} + rand_\eta(A'_{k-\eta}),$$

where $\eta = \lfloor k(0.1 + 0.9/e^x) \rfloor$ and x is the number of training episodes completed. This approach reduces both memory requirements and execution time by constraining the search space, and favors exploration at the beginning, when 100% of the actions are chosen randomly.

We illustrate the DRL agent’s structure in Fig. 2 and present the training process in Algorithm 1. The agent takes as input the sparse measurements snapshot M_t^{i-1} built up to the current

Algorithm 1: Training the RL agent

```

1 Input: Pre-trained MTRNet  $f()$ , environment  $E$ , the number
  of epochs  $n_e$ , the number of snapshots in training data  $n_s$ ,
  the number of historical snapshots used  $T$ ;
2 Initialize the neural network  $p()$ ;
3 for  $epoch \leftarrow 1$  to  $n_e$  do
4   Initialize  $E$ ;
5   for  $t \leftarrow T + 1$  to  $n_s$  do
6     while episode not finished do
7       Get current state  $s_i = [M_{t-T}, \dots, M_t^{i-1}]$ ,
        previously selected actions  $\mathbb{A}$ , and the ground
        truth snapshot  $F_t$  from  $E$  for the current
        iteration  $i$ ;
8       Generate subset  $A'$  from  $A$  and  $p(M_t^{i-1}, t, \mathbb{A})$ ;
9        $e = \{\}$ ;
10      for  $a$  in  $A'$  do
11        Get  $s_j$  by applying  $a$  to  $s_i$ ;
12         $e_j = MAE(f(s_j), F_t)$ ;
13        Push  $e_j$  to  $e$ ;
14      end
15       $a' = \arg \min_{a \in A'}(e)$ ;
16      Apply action  $a'$  to  $E$ ;
17       $Loss = \|p(M_t^{i-1}, t, \mathbb{A}) - a'\|_2^2$ ;
18      Use  $Loss$  to update  $p()$  by gradient descent;
19    end
20  end
21 end

```

iteration i , the list of cells selected so far, and the time t . From the sparse measurements matrix, the neural network first extracts feature maps using 2D convolution layers, with a sum pooling layer applied between them to summarize these feature maps without diluting the active features. This enables faster learning of the traffic patterns observed within a geographical window. The resulting features are then fed to a 2D convolution layer to reduce the size of the aggregate feature map, such that each feature’s magnitude is of the same scale as the time and previous actions. Two fully connected layers then process the aggregate feature map, along with the time and the previous actions, to predict the next action. Based on this value, the agent generates a set of potential candidates to be evaluated by MTRNet (Algorithm 1, lines 10–14), and the candidate yielding the smallest reconstruction error is chosen by the agent as next action (line 15). Finally, the neural network learns to improve its predictions by updating its weights based on gradients of the Mean Squared Error loss between the predicted (pseudo) action and the chosen candidate (lines 17–18). For any given snapshot, this process is repeated until the reconstruction error is less than a predefined reconstruction quality threshold.

In choosing this threshold we seek a trade-off between achieving low reconstruction errors and selecting a small number of cells. To find an appropriate quality threshold we examine the MAE reduction as a function of (random) cell sampling rates and observe that MAE decreases only marginally beyond 35% sampling rates. Hence we work with this value as the quality threshold.

When predicting the pseudo action, the agent only considers the current snapshot instead of the current state. This

is because in choosing actions one by one, only one value in a matrix changes from zero to the measurement between two consecutive states, which would make the collected value less important in the input, wrt. historical data. Although the previous experience is not seen in the pseudo action prediction, MTRNet evaluates an action given the full state as the input, whereby temporal correlations are captured.

B. Traffic Reconstruction from Sparse Measurements

We perform traffic measurement reconstruction using a supervised learning approach. The proposed Mobile Traffic Reconstruction neural Network (MTRNet) takes as input the sparse measurement matrices for the current (t -th) and previous T timestamps, i.e., $[M_{t-T}, \dots, M_t]$, and outputs the traffic consumption of the full map $\hat{F}_t$ at the current timestamp.

The model architecture is shown in Fig. 3. We draw inspiration from ZipNet [9], a mobile traffic super-resolution technique that infers fine-grained traffic consumption from coarse measurements. The original ZipNet was modified in order to reduce the size of the model and improve its inference speed. Specifically, since in our setting the size of sparse measurement matrices is the same as the size of the output, we discard upscaling blocks. Further ablation study allows us to reduce the number of layers of the model (depth) while retaining high accuracy. The resulting MTRNet comprises three key components:

- **Correlations capturing.** A 3D convolutional layer is used to capture spatiotemporal correlations between recent snapshots. The activation layer is a Leaky-ReLU (LReLU) that improves the model's non-linearity and its robustness to network initialization. LReLU is defined by:

$$\text{LReLU}(x) = \begin{cases} x, & x \geq 0 \\ \alpha x, & x < 0 \end{cases},$$

where α is a configurable slope value. A SumPooling layer is applied after LReLU to reduce the size of the feature maps and speed up the training.

- **Feature extraction.** Several 2D Convolutional layers followed by a LReLU extract high level abstract features from the geographical configuration of the measured traffic, thereby enhancing the representability of the model. Staggered skip connections link every two blocks, and a global skip connection links the input and output of this component. These skip connections hierarchically preserve different features extracted from previous layers, enabling feature reusability and stabilizing training and convergence.
- **Feature summarization.** To maximize the probability of capturing all relevant combinations of abstract features extracted by the previous component, the first two layers of the feature summarization block gradually increase the number of convolution filters (or kernels). Subsequently, the resulting feature maps are merged by a third and last convolution layer to provide a faithful reconstruction of the complete traffic snapshot, which represents the final output of our proposed MTRNet.

Fig. 3: MTRNet structure comprising correlations capturing, feature extraction & summarization functionality. Traffic snapshots reconstructed based on historical sparse measurements.

C. Policy Network

Since RL agents typically evaluate one action at a time, they are often impractical to deploy in settings where decisions have to be made within short deadlines. Therefore, we design a policy network that *predicts all optimal cell locations to be sampled*, $B_t = g(\mathbb{M}_t)$, *at once*. This policy network (g) takes historical sparse measurement snapshots as input and learns to predict binary selection matrices generated by the agent introduced in Sec. III-A at the end of each episode.

This neural network effectively solves a multi-label classification task where positive labels represent cells to be selected and negative labels indicate non-selected squares. The model architecture is similar to that of MTRNet, though here a final Sigmoid layer transforms the outputs into probability scores and we assign positive labels to those selection matrix elements where the probability is above 0.5, and negative to all others. As the cell selection frequencies are likely to be skewed, the probabilities are normalized to the range $[0, 1]$ before the binarization process.

We use a Binary Cross-entropy (BCE) loss function to train this neural model, which is defined by $BCE = -\frac{1}{N} \sum_{i=0}^N y_i \log(\hat{y}_i) + (1-y_i) \log(1-\hat{y}_i)$, where y_i is the value in the binary selection matrix, and $\hat{y}_i$ is the output probability score. N denotes the total number of elements (cells).

IV. EVALUATION

In this section we evaluate the performance of our Spider framework in terms of the number of cells sampled for mobile traffic visibility, and measurement reconstruction error.

We use a compute cluster comprising 20+ nodes, each equipped with 1-2 NVIDIA TITAN X GPUs (2280 cores) to train the neural models. We implement Spider in Python using the PyTorch library, and train on 2 GPU nodes in parallel for 2 epochs, spending 9.6 hours on each GPU. The number of historical snapshot T is 6. For all the models, we employ the Adam optimizer with learning rate $\lambda = 10^{-4}$ and use 128 as the batch size. For MTRNet, the number of previous actions is 20 and the number of 2D Convolutional layers is 5.

A. Dataset

For evaluation we adopt a real-world mobile traffic dataset collected by Telecom Italia in the city of Milan [10]. The city is divided into 100×100 squares of 0.055 km^2 . Mobile data traffic measurements were collected on aggregate at these locations every 10 minutes between 01/10/2013 and

Fig. 4: Number of cells selected for measurement collection by Spider and benchmarks, over one week (data representative for 16–22 Dec 2013). Average volume of traffic also plotted to highlight Spider’s ability to adapt to changing patterns.

TABLE I: Performance comparison between Spider and cell selection baselines, in terms of number of cells sampled and resulting NMAE following MTRNet-based reconstruction.

Interval	Random cell selection		Historical data-based selection		Spider	
	Count	NMAE	Count	NMAE	Count	NMAE
Peak	4,737	0.06	4,185	0.04	2,643	0.08
Off-peak	3,610	0.09	2,472	0.10	1,814	0.12
Weekdays	4,116	0.08	3,142	0.07	2,378	0.09
Weekend	3,838	0.08	3,033	0.07	2,237	0.09
Holiday	4,049	0.07	3,129	0.08	1,906	0.11
Overall	4,039	0.08	3,122	0.08	2,109	0.10

01/01/2014 (3 months). We use 40 days worth of data to train the models, and 20 days for testing. We normalize the traffic consumption and work with $\tilde{x} = \log(1+x)/\bar{x}$, where x is the traffic consumption and $\bar{x}$ is the mean of $\log(1+x)$.

B. Measurement Collection

We first examine Spider’s performance in terms of the average number of cells selected for measurement collection versus (1) a random selection strategy that chooses randomly which cells to sample, with their number being the average at any time of the day during different days of the week that yields reconstruction errors below the predefined threshold, as observed across the dataset used for training our agent; and (2) a strategy that chooses the most frequently selected cells based on past selection patterns. This frequency matrix is generated by averaging binary selection matrices obtained during training for the same time of the day and day of the week. To put things into perspective, we also compute the Normalised Mean Absolute Error (NMAE) obtained with our MTRNet when each of these strategies are employed.

We summarise the performances of Spider against the baselines considered in Table I, comparing NMAE and sampled cell count (rounded to the nearest integer) at peak and off-peak hours during a day, respectively during weekdays, weekends, and public holidays. Overall, **Spider samples 48.0% fewer cell than the random selection strategy and 32.4% fewer than the selection approach based on historical data**, at a negligible cost in terms of additional reconstruction error introduced relative to the average volume of traffic (NMAE).

We illustrate Spider’s behavior and that of the benchmarks considered over an entire week (Monday to Sunday) using traffic from the testing set (red line) in Fig. 4. Observe that Spider captures accurately the changes in traffic demand and samples cells for measurement collection accordingly.

Fig. 5: Number of cells sampled by Spider and the benchmarks considered, during Christmas holidays.

Fig. 6: Cell selection frequency. From left to right: historical data-based selection in off-peak, peak; Spider off-peak, peak.

While the historical data-based approach is able to distinguish between different times of the day and days of the week, it clearly over-samples, thus incurring higher measurement collection overhead. This is even more noticeable when cells are sampled randomly and with the only goal of meeting the quality threshold.

The superior performance of Spider is further emphasized in Fig. 5 where our framework is applied for measurement collection during the Christmas holidays, when traffic demand decreases below typical daily averages. Spider is able to adapt to previously unseen traffic patterns and distinguishes between holidays taking place on weekdays, and normal weekdays. Precisely, it selects 19.8% fewer cells on this occasion. In contrast, random and historical data-based selection approaches fail to adapt to such circumstances.

We delve deeper into which cells are selected by Spider, showing in Fig. 6 (right) their selection frequency at peak (7AM–7PM) and off-peak (7PM–7AM) times during weekdays, juxtaposed with the behavior of the historical data-based approach (left). Observe that Spider focuses more on the city center where the traffic demand is relatively larger, expanding sampling coverage at peak time.

Finally, we note that directly predicting a selection matrix by Spider’s policy network takes 7 milliseconds on average, whereas predicting the next best action by the RL agent would have taken 2 seconds. Thus **our approach is 284 times faster and suitable for operational settings**.

Fig. 7: Comparison between MTRNet and alternatives for traffic reconstruction in terms of MAE wrt. ground truth. Inference time in millisecond (ms) is shown above the bars.

Fig. 8: Example normalized traffic volume of snapshot predicted by CS, STCS, KNN and MTRNet, and respectively the ground truth (from left to right)

C. Traffic Reconstruction

Next we examine the performance of our MTRNet in terms of reconstruction quality, versus that of KNN, compressive sensing (CS) [6] and spatio-temporal compressive sensing (STCS) [7] alternatives. With KNN, traffic at unsampled locations is approximated by the K nearest sensed values with a weight of their distance. CS exploits and iteratively solves ℓ_1 minimization problems to reconstruct a sparse snapshot. STCS incorporates global spatio-temporal properties into CS.

Fig. 7 shows the average MAE attained by each approach when the sampling rate varies between 10–70%. MTRNet outperforms the benchmarks considered, achieving the smallest MAE. The performance gap is wider when the random sampling rate is small. Specifically, **MTRNet reduces the MAE of KNN, CS and STCS by 52.4%, 67.43% and respectively 66.8% when the sampling rate is 10%. Further, MTRNet reduces inference time/snapshot by up to 835×.**

Examining the predictions made by all algorithms when traffic is sampled randomly with 35% rate, Fig. 8 demonstrates that MTRNet achieves remarkable fidelity in reconstructing traffic snapshots. In contrast, CS and STCS over-estimate the traffic volume, while KNN is unable to accurately reconstruct traffic in areas with lower consumption.

V. RELATED WORK

Data reconstruction in urban tasks: Rana et al. [11] propose an urban noise mapping system where CS is employed to extrapolate incomplete and random samples obtained via crowdsourcing. He and Shin [12] propose an incentive mechanism based on Bayesian Compressive Crowdsensing to recover an urban signal map. Sparsity Regularized Matrix Factorization (SRMF) is used in [7] for traffic data interpolation. However, real-world traffic matrices rarely meet the conditions required by CS-like algorithms (e.g., Gaussian-like distribution, low rank, etc.). Zhang et al. [9] formulate traffic reconstruction as a super-resolution task and propose a generative model to

solve this task. Our MTRNet applies similar principles, but does not upscale and adapts to changing sampling strategies. **Sparse Mobile Crowdsensing (SMCS):** SMCS approaches select urban subareas via a policy, and recruit users covering these subareas to upload sensing data using mobile devices, subsequently inferring the missing data [13]. Compared to infrastructure-level data that is collected from dedicated probes, such user-level data collection paradigm has given rise to concerns including data unavailability, privacy issues and inefficiency. Wang et al. use Query by Committee (QBC) for cell selection and SRMF for data inference [14]. SMCS is later viewed as a RL process, where deciding what cell to sense (actions) is informed by data already collected (states) [15]. A Deep Q-Network (DQN) performs cell selection, which is unsuitable for large action spaces.

VI. CONCLUSION

In this paper we introduced Spider, a mobile traffic measurement collection and reconstruction framework that harnesses deep (reinforcement) learning to reduce measurement collection overhead and infer mobile traffic consumption at city-scale with high accuracy. We overcame the challenges of handling large action spaces in the cell sampling process and designed a purpose-built neural model that exploits spatiotemporal patterns characteristic to mobile traffic, to reconstruct network-wide traffic snapshots. We demonstrated Spider reduces cell sampling overhead by up to 48% while significantly outperforming existing measurement interpolation methods.

ACKNOWLEDGEMENTS

This work was supported by Cisco through the Cisco University Research Program Fund (Grant no. 2019-197006).

REFERENCES

- [1] “Ericsson Mobility Report,” [Online; accessed 1-April-2021], Nov 2020.
- [2] C. Zhang and P. Patras, “Long-term mobile traffic forecasting using deep spatio-temporal neural networks,” in *MobiHoc*, 2018.
- [3] D. Naboulsi et al., “Large-scale mobile traffic analysis: a survey,” *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 124–161, 2015.
- [4] X. Cheng et al., “Mobile big data: The fuel for data-driven wireless,” *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1489–1516, 2017.
- [5] S. Wu et al., “Dynamic traffic prediction with adaptive sampling for 5G HetNet IoT applications,” *Wireless Comms & Mob. Comp.*, 2019.
- [6] D. Jiang et al., “A compressive sensing-based approach to end-to-end network traffic reconstruction,” *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 1, pp. 507–519, 2020.
- [7] M. Roughan et al., “Spatio-temporal compressive sensing and internet traffic matrices,” *IEEE/ACM Trans. Netw.*, vol. 20, no. 3, 2011.
- [8] S. Ma et al., “Service demand prediction with incomplete historical data,” in *IEEE ICDCS*, 2019.
- [9] C. Zhang et al., “ZipNet-GAN: Inferring fine-grained mobile traffic patterns via a generative adversarial neural network,” in *CoNEXT*, 2017.
- [10] G. Barlacchi et al., “A multi-source dataset of urban life in the city of Milan and the Province of Trentino,” *Scientific data*, vol. 2, no. 1, 2015.
- [11] R. K. Rana et al., “Ear-phone: an end-to-end participatory urban noise mapping system,” in *IPSN*, 2010.
- [12] S. He and K. G. Shin, “Steering crowdsourced signal map construction via Bayesian compressive sensing,” in *INFOCOM*, 2018.
- [13] L. Wang et al., “Sparse mobile crowdsensing: challenges and opportunities,” *IEEE Comm. Mag.*, vol. 54, no. 7, pp. 161–167, 2016.
- [14] —, “Cell selection with deep reinforcement learning in sparse mobile crowdsensing,” in *ICDCS*, 2018.
- [15] W. Liu et al., “Multi-dimensional urban sensing in sparse mobile crowdsensing,” *IEEE Access*, vol. 7, pp. 82 066–82 079, 2019.